

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
17.09.2003 Bulletin 2003/38

(51) Int Cl.7: **G06F 9/46**

(21) Application number: **02259029.3**

(22) Date of filing: **31.12.2002**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR
IE IT LI LU MC NL PT SE SI SK TR
 Designated Extension States:
AL LT LV MK RO

(72) Inventors:
 • **Large, Andrew R.**
La Selva Beach, CA 95076 (US)
 • **Prabandham, Harish**
San Jose, CA 95117 (US)

(30) Priority: **14.03.2002 US 364480 P**
14.03.2002 US 364862 P
19.09.2002 US 251084

(74) Representative: **Frost, Alex John et al**
Boult Wade Tennant,
Verulam Gardens
70 Gray's Inn Road
London WC1X 8BT (GB)

(71) Applicant: **Openwave Systems Inc.**
Redwood City, CA 94063 (US)

(54) **Method and apparatus for requesting and performing batched operations for web services**

(57) The Web Services Development Kit (WSDK) comprises a set of build-time tools and run-time components designed to support the development of web services and the interface to those services. WSDK provides a framework for the development of web services that isolates developers from the implementation details of the markup and transport layers, by separating the log-

ical aspects of the service development from the physical aspects. A "call" generated by a client contains a batch of operations to be executed in the context of a single request/reply message exchange. A markup library translates logical structure to a specific markup language. Additional markup languages are supported by adding additional modules to the markup library.




FIG. 1

Description

[0001] This application claims the benefit of U.S. Provisional patent application no. 60/364,480 filed on March 14, 2002 and U.S. Provisional patent application no. 60/364,862 filed on March 14, 2002, each of which is incorporated herein by reference.

FIELD OF THE INVENTION

[0002] At least one embodiment of the present invention pertains to tools and components for the development of web services and, more particularly, to a method and apparatus for requesting and performing batched operations for Web Services.

BACKGROUND

[0003] "Web services" is an industry term that covers a collection of standards and proposals that define how distributed computing should be implemented in a web-based network environment, such as the World Wide Web. Web services is an evolving area of technology that has a broad range of potential applications but is immature in implementation. There are several general attributes that describe the essence of web services.

[0004] The primary paradigm for web services is that of a client/server relationship. A client sends a request to a server for the execution of some task. The server executes the task and sends an appropriate reply back to the client.

[0005] Interoperability is a main goal of current web services efforts. Applying lessons from past mistakes, the industry is focusing its efforts on ensuring that information communicated between components of the system will be interpreted identically by all of the parties involved. The focus on "on the wire" representation has allowed for the development of very different (and competing) technologies for constructing web service implementations. Two of the major contenders at present include Microsoft, with its .NET architecture that leverages its traditional COM technology, and IBM, with its WebSphere architecture that leverages such Java technologies as Java 2 Enterprise Edition (J2EE).

[0006] Communication between client and server on the World Wide Web is typically accomplished via Hypertext Transport Protocol (HTTP). This is one of the factors that distinguish web services from other client/server approaches, such as remote procedure call (RPC) or common object request broker architecture (CORBA). Web service interactions also differ from those of typical Web browser interactions. In the typical browser interaction, the client sends a request for content from the server which will be rendered for display to an end user. In the web services model, the client sends a request for a task (or operation) to be executed by the server, and the reply contains the results of executing that operation.

[0007] Underlying most aspects of web services is the extensible markup language (XML). Most situations that call for representation of data employ XML as the base data structuring language. Most web services protocols are defined as a set of XML elements and conventions. This is another major factor that distinguishes web services from other client/server approaches. For example, CORBA employs one language (IDL) for defining interfaces and a different protocol (IIOP) for communicating between client and server. In the web services context, both the interface definition language and the client/server communication protocol are represented in XML.

[0008] There are two problems with XML, however, that probably originate from its original focus on supporting text processing. First, XML can be quite difficult and expensive to parse, due to such features as user-defined entities, processing instructions, etc., most of which are not needed for most inter-application communication. Second, as a text-based protocol, XML does not support the transport of raw binary information. Instead, it must be encoded into text for transmission (e.g., Base64 or similar encoding mechanisms).

[0009] There are essentially two approaches typically used for solving the first problem. The first approach is to define protocols that are based on the important core features of XML, but specifically forbid the use of particularly complex XML features (such as processing instructions). A recent example of this approach is the Simple Object Access Protocol (SOAP). See "Simple Object Access Protocol (SOAP) 1.1," W3C Note, May 8, 2000. SOAP is an XML-based protocol that provides basic rules for the format of request and reply messages exchanged between client and server during operation execution. The second approach is to define a separate "logical" document model and then provide information to "map" that logical model onto the specific XML representation. A recent example of this approach is the Web Services Description Language (WSDL). See "Web Services Description Language (WSDL) 1.1," W3C Note, March 15, 2001. WSDL is an XML-based interface definition language that is used to define the operations supported by a server and how requests and replies for those operations are represented in the messages exchanged between client and server.

[0010] SOAP defines four basic things:

- 1) A subset of XML grammar that is allowed in a SOAP message (basically, it disallows certain complex XML

constructs that are difficult to parse and are not applicable to a remote method invocation paradigm);

2) A basic document structure for SOAP messages that include:

- 5 • a top level SOAP "Envelope" element within which all SOAP-related XML data must reside,
- a SOAP "Header" element which is intended to contain all information pertaining to routing and handling of the SOAP document, and
- 10 • a SOAP "Body" element which is intended to contain all information pertaining to execution of the specific request (or reply);

3) How "faults" (document processing or task execution errors) should be represented within the document; and

- 15 4) A suggested approach for representing typed data in SOAP documents (e.g., mappings from low-level host programming language variable types to XML).

[0011] WSDL defines a grammar for describing web service interfaces. A WSDL interface description consists of the following five top-level elements:

- 20 1) A single *types* element. Within this element, you define the logical types that you will be referencing later in the specification. Types may be *simple* (refer to single elements) or *complex* (combinations of elements, e.g. struct or array). Different typing systems may be used here, but the most common one employed is XML Schema;
- 25 2) A set of *message* elements. A *message* element corresponds to a single logical "message" that will be sent from client to server or vice versa. Messages consist of "parts" - each part has a type (defined within the *types* element);
- 30 3) A set of *portType* elements. A *portType* element corresponds to single logical service definition. Port types consist of a set of operation definitions. An operation has a name and the set of messages involved in requesting execution of that operation (typically a *request* message, a *reply* message, and a *fault* message);
- 35 4) A set of *binding* elements. Up to this point, everything covered in the WSDL document represents a "logical" view of the service. The purpose of the binding elements is to define how the logical messages are represented in a particular markup language/transport. Support theoretically exists for several XML-based transports, but the most common binding is to SOAP/HTTP. In that case, the binding involves mapping each part of each message to the appropriate section of the SOAP document and identifying the approach that will be used for encoding the logically-defined data types in SOAP (typically the SOAP encoding rules are used here); and
- 40 5) A *service* element. Where the binding element is used to map the logical representation to a markup/transport, the *service* element is used to identify the physical locations that support a particular portType/binding pair. For the typical SOAP-HTTP binding, this would normally consist of uniform resource locators (URLs) (including host and port).

45 **[0012]** A full WSDL specification thus consists of logical information (types, messages and portTypes), physical information (bindings), and location information (services). Since this information is typically needed by clients at different points in the development, WSDL specifications are often produced in fragments that are assembled together via an analog of the C pre-processor "#include" directive.

50 **[0013]** Aside from SOAP and WSDL, there are other initiatives that are often considered web services "building blocks", such as XML schema and universal description, discovery, and integration (UDDI) protocol. XML schema provides a set of rules for defining logical data types in XML and is the preferred language for defining types in WSDL. UDDI is an XML-based protocol for obtaining information about the services provided by a particular organization. Typically, UDDI is used to retrieve the complete WSDL specifications (including location information) for a set of operations.

55 **[0014]** Combined, SOAP and WSDL perform three basic tasks when mapping information exchange onto a structured document:

- 1) Define a logical model for the interchange (the "parse tree" view of the problem). WSDL provides this.

2) Define a physical model for the interchange (the "on the wire" representation). SOAP provides this.

3) Provide a mapping between the logical and physical models. WSDL provides this.

[0015] Consider the following sequence of events, which describes a typical experience for a web services client developer:

- The developer first obtains the "logical" portion of WSDL specifications for operations of interest via an "out of band" method, such as email or downloading from a web site.
- The Developer imports the WSDL specifications into a web services toolkit and generates a native language (e.g., Java) client application programming interface (API) and writes code to that API.
- At runtime, a complete WSDL specification for the operations of interest is obtained via UDDI from a service "registry" at the deployed site. The additional binding and location information in the "complete" WSDL files is necessary for actual communication between client and server (though it does not change the client API originally generated).
- When runtime calls are made via the generated client API, they are translated into SOAP documents that are sent to the server via an HTTP POST. The server executes the requested operation and returns the results as a SOAP document via the HTTP POST reply. The returned document is parsed and used to populate the appropriate "output" parameters in the client API.

[0016] The two basic approaches represented by SOAP and WSDL have shortcomings. For example, SOAP does not define how one represents the tasks to be executed by the server, how the data associated with the request and reply is structured, relationships between request and reply messages, or real-world issues such as application versioning. These issues must be coordinated between client and server "out of band". WSDL addresses the "logical to physical" mapping problem by individually specifying how different elements in the logical model map to their counterparts in the physical model. Consequently, a consumer of such an interface must parse and understand such mapping instructions, which is not a trivial task. In addition, SOAP and WSDL are both XML-based. Consequently, without either changing the XML specification or manipulating the underlying transport (e.g., HTTP) protocol, one cannot pass binary data without encoding it as text. Other problems are associated with SOAP and WSDL, as described below.

[0017] Further, additional facilities are needed for effective implementation of typical web services, which are not provided by known current specifications. These facilities include user identification and authentication, access control, version management, transaction management, and ad hoc (dynamic) "batching" (sending several logical operation requests together in a single physical message).

SUMMARY OF THE INVENTION

[0018] The present invention includes a method and a corresponding apparatus, where the method includes dynamically determining in a client a plurality of batched operations to be performed by a server for the client during execution of a web based network application. This is accomplished by selecting a portion of an operation structure previously defined using a structured language. The method further includes operating the client to specify the plurality of batched operations in a request, and sending the request from the client to the server to request the server to perform the plurality of operations in a batched mode.

[0019] Other features of the present invention will be apparent from the accompanying drawings and from the detailed description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] One or more embodiments of the present invention are illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

[0021] Figure 1 illustrates a high-level view of the WSDK architecture;

[0022] Figure 2 shows a process performed by a service developer during a WSDK-based development project;

[0023] Figure 3 shows an example of a network environment in which WSDK can be implemented;

[0024] Figure 4A shows components of the WSDK Server Framework;

[0025] Figure 4B shows components of the WSDK Client API Framework;

[0026] Figure 5 is a high-level block diagram of the WSDK Library;

[0027] Figure 6 shows components of the input module of the WSDK Library;

[0028] Figure 7 shows an overall process for a document exchange between a client and a server using the WSDK APIs;

[0029] Figure 8 shows a process performed by the client during the document exchange;

5 [0030] Figure 9 shows a process performed by the server during the document exchange; and

[0031] Figure 10 is a block diagram showing an abstraction of a processing system that can be used to implement the WSDK and related techniques.

DETAILED DESCRIPTION

10

[0032] A method and apparatus for requesting and performing batched operations for Web Services are described. Note that in this description, references to "one embodiment" or "an embodiment" mean that the feature being referred to is included in at least one embodiment of the present invention. Further, separate references to "one embodiment" or "an embodiment" in this description do not necessarily refer to the same embodiment; however, such embodiments are also not mutually exclusive unless so stated, and except as will be readily apparent to those skilled in the art from the description. For example, a feature, structure, act, etc. described in one embodiment may also be included in other embodiments. Thus, the present invention can include a variety of combinations and/or integrations of the embodiments described herein.

15

I. Overview

20

[0033] As described in greater detail below, the present invention introduces a set of build-time tools and run-time components designed to support the development of web services and the interface to those services. These tools and components are referred to collectively as Web Services Development Kit (WSDK). WSDK provides a framework for the development of web services that isolates developers from the implementation details of the markup and transport layers. In particular, WSDK separates the logical aspects of the service development from the physical aspects (markup and transport) so that service developers can focus on business logic and platform developers can take care of mapping the logical model into particular markup languages (or versions of particular languages). Note that for purposes of this document, for ease of reference, different versions or variations of a given markup language shall be treated as different markup languages. WSDK provides additional support for real-world application development and deployment issues not addressed in the standard web services arena, such as batching, transaction management, and versioning, etc.

25

[0034] WSDK provides the above features in a form suitable for the extreme performance and scalability requirements of a carrier-class environment. For example, in certain embodiments of the invention: all server-side code is written in low-overhead C and C++; there is minimal use of dynamic memory allocation and string/data copying; hand-rolled "no copy" parsers are used with little or no dynamic memory allocation; and efficient code generated at build-time is used rather than run-time data dictionary queries.

35

[0035] There are three main areas where WSDK differs conceptually from the service model defined by WSDL. First, the concept of "service" differs between WSDL and WSDK. In WSDL, a service may be viewed as a set of network end-points (host, port, and URI) that support a specific collection of "bound" operations. WSDK, on the other hand, distinguishes between "server" (a run-time software unit that handles requests on an end-point) and "service" (a well-defined collection of operations more akin to the WSDL "portType"). A WSDK server may host multiple services. In the remainder of this document, "service" will be used to mean a top-level aggregation of operations and related data types and fault codes.

40

[0036] Second, WSDL defines a direct relationship between operations and messages. Essentially, the entire content of a message is defined by the logical content and binding method of a single operation. WSDK, on the other hand, introduces an intermediate runtime-only entity known as a call. A call contains a batch (one or more) of operations to be executed in the context of a single request/reply message exchange. Batching is an optional runtime decision made by the client and is only available via the client APIs provided with WSDK. A batch of operations is defined dynamically by a client, based on a previously defined (static) operation structure definition. Thus, a client may dynamically select particular portions of an operation structure definition to incorporate into a message (a batched request), according to its current needs and/or the current state of application execution. Note that messages generated from a call containing only one operation are identical to those defined via WSDL.

45

50

[0037] Further, in WSDL the relationship between the logical operation request/reply structure and a particular markup language representation is specified via a declarative per-operation syntax. WSDK, on the other hand, uses an entirely different mechanism, employing a markup library that can dynamically translates logical structure to any supported markup language, without using any explicitly-declared binding of logical structure to markup language encoding. Additional markup languages are supported by adding additional modules to the markup library. Where

55

explicit (e.g., WSDL) bindings are required, the hard-coded internal rules are simply translated into WSDL binding statements.

II. WSDK Architecture

[0038] To facilitate description of the present invention, various components are described herein as being implemented in a particular programming language, such as C, C++, or Java. It will be recognized that such references to particular programming languages refer only to certain embodiments of the invention, and that in other embodiments, the described components may be implemented in other programming languages and/or by using hardware-based techniques.

[0039] Figure 1 depicts a high-level view of the WSDK architecture according to one embodiment of the invention. There are four types of components depicted in Figure 1: components of WSDK (indicated by thick solid lines); components related to WSDK but not part of it (indicated by thin solid lines); Operation Structure Definition (OSD) specifications; and generated code (depicted in dashed boxes).

[0040] The components of WSDK include a WSDK Code Generator¹, a WSDK Client API Framework (WCAF) ², a WSDK Server Framework (WSF) ³, and a WSDK Markup Library ⁴. The Code Generator ¹, henceforth referred to as "OSDGen", is a build-time generation tool that parses Operation Structure Definition (OSD) files and generates code for both service implementation and client interface.

[0041] The WSF ³ is, in one embodiment, a C++ server framework that manages the runtime operation dispatch process and all associated facilities.

[0042] The WCAF ² is, in certain embodiments, a C or Java client framework that provides a framework within which client APIs are exposed. As described further below, this framework includes an "OipcCall" data structure that is populated with "OipcOperation" objects. The WCAF ² also includes a set of functions to serialize the request portion of the call to a request document and to populate the reply portion of the call from a reply document. In general, these "call-level" functions are simply wrappers over the operation-specific code generated by OSDGen ¹.

[0043] The Markup Library ⁴ is, in certain embodiments, a C or Java markup library that translates logical markup API calls into physical markup (e.g., SOAP). This library ⁴ is used by both the WSF ³ and the WCAF ².

[0044] The illustrated components related to WSDK include: a Server Transport Framework ⁵; Operations, Administration, and Management (OA&M) facilities ⁶; and a Client Transport Library ⁷. Multiple implementations of these components may exist.

[0045] The server transport framework (STF) ⁵ is responsible for the basic runtime environment the server side. It may be, for example, a UNIX program that is designed to service "request" messages sent from a client over a TCP socket. The STF ⁵ would typically be capable of servicing multiple requests simultaneously (so it typically either has multiple threads or multiple processes) and may use HTTP as a "tunneling" protocol (e.g., requests are sent via an HTTP POST). The STF ⁵ would typically allow for the separate development and integration of modules that implement service business logic.

[0046] The WSF ³ is provided in the form of a library that is linked into a main program, i.e., the STF ⁵. The STF ⁵ "drives" the WSF ³ via two interfaces: a "startup initialization" interface, and a "request dispatch" interface. An example of the STF ⁵ for WSDK-based servers is the Apache web server. This is combined with a small Apache module that provides the "glue" between Apache and WSF ³.

[0047] Thus, the "server-side" can be viewed as having three levels of abstraction:

- 1) The "server" level (provided by the STF) - a "main" program.
- 2) The "service framework" level (provided by the WSF ³) - a "module" in the form of a library that is "plugged in" to the server main program.
- 3) The "service" level - a set of shared libraries (one for each service) that is "plugged in" to the WSF ³.

[0048] The OA&M facilities ⁶ are a set of configuration and logging facilities used by the WSF ³ via an abstract interface. Typically, the STF ⁵ is responsible for initializing appropriate instances of these facilities.

[0049] The Client Transport Library ⁷ is used by the client for communication with WSDK-based servers when the generated WCAF ² is employed. The Client Transport Library ⁷ provides basic TCP/IP connection management and client-side HTTP support.

[0050] Operation Structure Definition (OSD) Specifications ⁸ are input files that are provided by developers and are described further below.

[0051] The generated code shown as dashed-line boxes ⁹ in Figure 1 includes components generated by OSDGen ¹ and described in greater detail below.

[0052] The Service Implementations are the server-side code (e.g., C++) generated by OSDGen ¹. They include a combination of type declarations, service initialization, operation serialization and validation, and skeletons for opera-

tion business logic.

[0053] The Service Client APIs 9A are the client-side code (e.g., C or Java) that corresponds to individual operations. This code includes a combination of type declarations and operation serialization logic.

[0054] The Service WSDL Specifications 9B are the WSDL specifications (logical and binding descriptions only) that correspond to the operations of the service. They are provided as an alternative to using the client APIs. Note that WSDL location information (e.g., host and port) is not provided by WSDK, since it depends upon information known only at deployment. Clients that intend to use WSDL specifications must obtain the location information via another facility.

[0055] Referring now to Figure 2, a WSDK-based development project will normally include the following process performed by the service developer:

- Compose the OSD specification for the service (block 201), i.e., construct the OSD file(s) for the service being defined, including the specifications for each of the operations the service will provide. OSD and OSDGen 1 are described further below.
- Run OSDGen 1 (block 202) to create the server-side operation context classes and the skeletons for the server-side service and operation classes.
- Complete the implementation of the generated server skeleton classes and build the service into a shared library (block 203). When this library is placed in a designated directory, the service will automatically be loaded and initialized during WSF startup initialization. This stage includes developing the set of utility classes that will be needed by the operation during execution. Typically, this might include code for interfacing with database systems, authenticating users, etc. Where appropriate, implement the "Transaction Member" interface for participation in the message-defined transaction model.
- Run OSDGen 1 again (with different options) to generate the client-side interface artifacts (block 204), including: Logical and binding portions of WSDL specifications for the service. Fault message catalogs that can then be localized for display of human-readable fault messages; and client APIs.
- Implement the execute() methods for each of the operation classes (block 205).

II. Network Environment

[0056] Figure 3 shows an example of a network environment in which WSDK and other aspects of the present invention may be implemented. Note, however, that the invention can also be implemented within many other network configurations. As shown, a wireless client device 31 is connected to a mobile access gateway (MAG) 32 through a wireless network 33. The wireless client 31 device may be, for example, a cellular telephone, personal digital assistant (PDA), two-way pager, or other similar device. The MAG 32 is connected to a number of servers 34 through a wireline computer network ("wireline network") 35, which may be the Internet, for example. A second client device 36 is connected to the server directly through the wireline network. Client device 36 may be, for example, a conventional personal computer (PC), workstation, or other similar device.

[0057] The MAG 32 controls the connection to the wireless client device 31 (session management), retrieving content from the origin server 4 on behalf of the wireless client device 31 (proxying), and ensuring that the content is in a form that the wireless client device 31 can understand (translation). The MAG 32 may also perform many other functions, such as billing, enhanced services, and push content delivery.

[0058] In various other embodiments, the wireline network 35 could be a type of network other than the Internet, such as a local area network (LAN), a private wide area network (WAN), or a corporate intranet.

[0059] One or more wireless client devices 31 can be connected to one or more servers 34 on the wireline network 35 through the MAG 32. Likewise, one or more other client devices 36 can be connected to one or more servers 34 in this manner directly through the wireline network 35. Any of the clients and servers shown in Figure 2 may implement WSDK and other features of the invention described herein.

[0060] As noted, the invention can also be implemented within many other network environments. For example, the invention can be implemented advantageously in purely-wireline network environments, i.e., in network environments which do not included a wireless network or wireless (mobile) devices.

III. WSDK Server Framework (WSF)

[0061] Figure 4A shows the key classes/objects in the WSF 3 and the high-level relationships (control flow) between them. The classes include: OipsCallMgr, OipsCallResources, OipsServiceMgr, OipsCallCtxt, and OipsOpCtxt. OipsCallResources has a subclass OipsTxnCtxt, which has a subclass OipsTxnMember. OipsServiceMgr has a subclass OipsService, which has a subclass OipsOperations. OipsOpCtxt has subclasses OipsOpSummary and OipsStatus.

[0062] Following each class name in Figure 4A are three additional items of information (in the following order):

- The number of subclasses of the given class that may exist. If the number is 0, the class may not be subclassed. For both subclass and object cardinality, "n" and "m" mean "any number".

5 • The number of instances (objects) of that class, relative to the enclosing object. If a class is an abstract parent class (listed in *italics*), then this refers to the number of instances of each subclass. Otherwise, it is the number of objects of the listed class.

- The lifetime of objects of that class.

10 **[0063]** There are three possible lifetimes of an object: server, call, and txn. Objects with "server" lifetime are created and initialized during server initialization and exist for the lifetime of the server. Objects with "call" lifetime are created and initialized when a new request message is received by the call manager and exist until the reply message has been generated. Objects with "txn" lifetime are created and initialized at the beginning of a logical transaction. Depending upon the error handling method identified in the request message, a transaction will either span an entire call or
15 just the execution of a single operation (for calls that contain only one operation request, this distinction is moot).

[0064] The roles of the key classes/objects of the WSF 3 will now be described. The OipsCallMgr object coordinates the call dispatch process. This object is responsible for initializing the call context and resources, for dispatching the operation executions, and for all major error handling and transaction management logic. The OipsCallMgr is a concrete class (no subclassing) with a single instance that exists for the lifetime of the server. The functionality of this class is described further below.

20 **[0065]** The OipsCallResources object is a holder for call-related resources that do not directly correspond to items in the request or reply messages.

[0066] The OipsTxnCtxt object contains a list of members (instances of subclasses of OipsTxnMember) of the current transaction and delegates commit() and rollback() methods to those objects.

25 **[0067]** The OipsCallCtxt object contains the parsed request message and partially generated reply message for the current call and provides methods for traversing the request and populating the reply. This class is the server-side equivalent to "OipcCall" on the client side (see description of WCAF, below). Hence, OipsCallCtxt includes:

- requestDoc (OipmInputDoc) - the parsed request received from the client
- 30 • replyDoc (OipmOutputDoc) - the reply document sent (or will send, when it is completed) to the client.
- errorMode (enum OipcErrorMode) -- the indication of the error mode to be used (e.g., "stop", "continue", or "rollback") when executing the operations in a request. In "stop" mode, execution of the operations for the request is stopped in response to an error occurring during execution of the operations. In "continue" mode, execution of the operations for the request continues in response to an error occurring during execution of the operations. In "roll-
35 back" mode, execution of the operations for the request is stopped in response to an error occurring during execution of the operations, and the state that existed immediately prior to execution of the first operation for the request is resumed.

40 **[0068]** The OipsServiceMgr object is responsible for loading service implementations (during server initialization) and for providing access to service operations (during request processing).

[0069] The OipsService class is an abstract parent class. Each implemented service will have one subclass of OipsService and one instance of that subclass. Service classes are primarily responsible for initializing operations (during server initialization) and for providing access to operations (during request processing). Service class skeletons are generated by OSDGen 1. In many cases, the generated code is sufficient.

45 **[0070]** The OipsOperation class is an abstract parent class. Each implemented operation will have one subclass of OipsOperation and one instance of that subclass. Operation class skeletons are generated by OSDGen 1 and contain an execute() method that is invoked by the OipsCallMgr object to execute a particular operation. Many subclasses of the OipsOperation class may exist. For each subclass, there is exactly one instance that exists for the lifetime of the server.

50 **[0071]** The OipsOpCtxt class is an abstract parent class that contains the "request" and "reply" fragments for each operation found in the request. OipsOpCtxt is equivalent to "OipcOperation" on the client side. Each implemented operation will have one subclass of OipsOpCtxt. Instances of these subclasses are created as necessary (by the OipsCallMgr object via delegation through the OipsServiceMgr object) to hold the data contained in the current request. Operation context classes are fully generated by OSDGen 1 and do not require any manual coding. OipsOpCtxt contains the following standard methods:

- parseRequest() - similar to "docToReply" function of OipcOperation. This populates the "request" portion of the operation context from data parsed from the requestDoc of the call context.

- writeReply() - similar to "requestToDoc" function of OipcOperation. This writes the "reply" data from the operation context into the replyDoc of the call context.

[0072] OipsOpSummary objects contain the set of information that will be logged in the operation summary logs for the execution of the associated operation. There is a standard set of information provided automatically by WSF. Individual operations can augment this information by adding additional parameters (with names of the form "x-name").

[0073] OipsStatus objects are used for reporting application faults. A standalone OipsStatus population function is automatically generated by OSDGen 1 for each faultdef defined for a service.

[0074] The objects and processes used on the server side are described further below.

IV. WSDK Client API Framework (WCAF)

[0075] The WCAF 2 will now be described with reference to Figure 4B. In the illustrated embodiment, the WCAF 2 includes two data structures 41, named OipcOperation and OipcCall, and a set of functions 42, named oipclnitCall, oipcAddOperation, oipcGenerateRequest, and oipcSetReply.

[0076] The OipcOperation data structure represents an instance of a particular operation. This structure is utilized in a manner similar to a C++ abstract class -- it defines a set of standard attributes and function pointers. Actual operations have their own unique structure types (but may be "casted" to OipcOperation). The OipcOperation structure contains the following:

- requestToDoc(OipcCall*, OipcOperation*)- write the request data from the operation into the "output document" of the OipcCall.
- docToReply(OipcCall*, OipcOperation*) - read from the "input document" of the OipcCall and write to the "reply" portion of the operation.
- cleanup(OipcOperation*)- clean up any memory allocated by "docToReply".

[0077] Each "derived" OipcOperation structure also contains a "request" structure and a "reply" structure.

[0078] The OipcCall structure contains the contents of an entire message exchange between client and server, i.e.:

- requestHeader (struct OipcRequestHeader) - a simple structure containing standardized request message header information (mostly client identification information used for logging).
- replyHeader (struct OipcReplyHeader) - a simple structure containing standardized reply message header information (mostly server identification information).
- ops (OipcOperation*) - an array of "OipcOperation" structures.
- markupStyle (enum OipmMarkupStyle) - the markup "style" used for this call (e.g., Lite, SOAP). A markup style corresponds to a markup language. It includes an enumerated type and an associated markup implementation (a software module) that handles conversion between the "generic" data format and a markup language-specific data format.
- errorMode (enum OipcErrorMode) - the error mode the server should use when executing the request (continue, stop, or rollback).
- requestDoc (OipmOutputDoc) - the request document created when the "requestToDoc" functions for all the operations in "ops" are called.
- replyDoc (OipmInputDoc) - the reply document parsed from the reply buffer received from the server.

[0079] The functions included in the WCAF 2 are oipclnitCall, oipcAddOperation, oipcGenerateRequest, and oipcSetReply. These functions are defined as follows:

- oipclnitCall(OipcCall*, OipmMarkupStyle, OipcErrorMode)- initialize a new call
- oipcAddOperation(OipcCall*, OipcOperation*)- Add an operation to this call (the "request" portion of the operation structure should be filled out before calling this function).
- oipcGenerateRequest(OipcCall*, char**, unsigned*) - Finalize the request message and return it as a text buffer for submitting to the server (calls the "requestToDoc" functions for each operation in the call).
- oipcSetReply(OipcCall*, char*, unsigned) - Parse the given reply buffer received from the server and use it to populate the reply part of each of the call's operations (calls the "replyToDoc" functions for each operation in the call).

[0080] As can be seen from the foregoing description, the WCAF 2 has components that mirror the functionality of the server side, but with some important differences. On the client side, a call is created from scratch. Separately,

operations are created from scratch, their request portions are populated, and they are added to the call. The call is then serialized into a request message and sent to the server. The resulting reply message is parsed and used to populate the "reply" portions of each of the operations.

[0081] In contrast, on the server side, a call is initialized with the incoming request document. No operation structures are contained within the call structure. Instead, operation structures are created as the request document is parsed and the reply document is written.

V. WSDK Library

[0082] The WSDK Library 4 (also called the "markup engine" or "markup library") has two modules: an input module 51 and an output module 52, as shown in Figure 5. Before examining these two modules in detail, however, note first that there are three principals in the markup space: 1) the "core" markup library; 2) markup implementations (described below); and 3) users of the markup library. Between these three principals, there are four distinct visibility relationships, that apply to both the input module 51 and the output module 52:

- The core markup library exports public data structures ("core-public-data") for use by both users of the markup library and markup implementations.
- The core markup library exports a public interface ("core-public-func") for use by users of the markup library.
- The core markup library exports a private interface ("core-private") for use by markup implementations ("core-private").
- Markup implementations export private interfaces ("markup-private") for use by the core markup library ("markup-private").

[0083] The above topology applies to both the input module 51 and the output module 52 of the WSDK Library 4.

A. Input Module

[0084] In accordance with one embodiment of the invention, the input module 51 includes the following elements:

- A single complex (core-public-data) data structure, called OipmlInputDoc, used by both markup library users and markup implementations.
- A set of (core-public-func) functions for "reading" the data structure, used by markup library users.
- A set of (core-private) functions for "writing" the data structure, used by markup implementations.
- A single (core-public-func) function for "creating" the data structure, used by markup library users. For each markup implementation, there is a single (markup-private) function that is invoked from the core-public-func function. This markup-private function uses the core-private functions to construct the OipmlInputDoc.
- A single (core-public-func) function for "destroying" the data structure, used by markup library users.

[0085] Thus, the input module 51 is illustrated conceptually in Figure 6, according to one embodiment. As shown, the input module 51 includes the data structure 61, called OipmlInputDoc (core-public-data), a submodule 62 (core-public-func) that handles the lifecycle and navigation of the input parse tree, a submodule 64 (core-private) that handles construction of the input parse tree, and a set of input markup implementations 63 (markup-private).

[0086] The OipmlInputDoc data structure 61 is a complex data structure that is used by both markup library users and markup implementations 63. It includes the following fields: 1) a "markup style" (i.e., an enumerated type defined as part of (core-public-data) and an associated markup implementation 63 that handles converting between the "generic" data format and a markup language-specific data format); 2) a reference to the raw text buffer that was used to construct the input document (this buffer is modified as the document is parsed); 3) an array of "entries"; 4) an array of entry "entry arrays"; and 5) a "current group" context. The "markup style" indicates the markup style that is in effect for the document.

[0087] An "entry" represents an entity found while parsing the document. There are two basic types of entries (indicated with a Boolean flag): an "item" entry and a "group" entry. An "item" entry has a single value associated with it. A "group" entry, which has a list of "child" entry arrays associated with it (see below). The parent-child relationship between groups and their sub-entries directly corresponds to the typical hierarchy one observes from structured documents. Consider the following example:

```
<foo>
  <bar1>value1</bar1>
  <bar2>value2</bar2>
</foo>
```

[0088] In this case, the group "foo" has two children (bar1 and bar2).

[0089] An entry array provides a mechanism for aggregating multiple entries with the same name. Entry arrays have a "name" and one or more child entries (the child entries may be items or groups). The parent-child relationship between entry arrays and entries is not typically apparent from an examination of original document -- it is an artifact of the implementation to support asking for entries by name and index.

[0090] Consider the following document fragment:

```
<foo>
  <bar>bar value 1</bar>
  <bar>bar value 2</bar>
  <fubar>fubar value</fubar>
</foo>
```

[0091] This write in would be represented in the parse tree as follows:

- * Entry array with name "foo"
 - * Group entry (index 0)
 - * Entry array (index 0) with name "bar"
 - * Item entry (index 0) with value "bar value 1"
 - * Item entry (index 1) with value "bar value 2"
 - * Entry array (index 1) with name "fubar"
 - * Item entry at index 0 with value "fubar value"

[0092] As noted above, the input module 51 of the markup library 4 includes a submodule 62 (core-public-func) that handles the lifecycle and navigation of the input parse tree, and a set of markup implementations 63. A set of key functions for operating on the input document form at least a portion of submodule 62 and markup implementations 63. These functions, according to one embodiment, are named: oipmInitInputDoc, oipmNumEntries, oipmOpenGroup, oipmGet* (where "*" represents one of several possible suffixes), and oipmCloseGroup, and are characterized as follows:

- oipmInitInputDoc(IN inputDoc, IN buffer, IN markupStyle) - this function inputs a text buffer and a markup style (how the text buffer should be interpreted). The function is called with an empty parse tree (no entries or entry arrays). When it returns, the parse tree is fully populated. All of the other input functions operate on the parse tree, not on the original text buffer. Documents have a single top-level group and the "current group" is set to this top-level group.

[0093] Note that documents written in any markup language (within reasonable bounds) can be converted into a standardized parse tree -- thus document navigation is fully insulated against differences in markup languages. Note also that for purposes of this document, for ease of reference, different versions or variations of a given markup language shall be treated as different markup languages.

- oipmNumEntries(IN inputDoc, IN name, OUT count) - This function looks for an entry array within the current group with the given name and then returns the number of child entries for that array. If no array is found with the given name or it contains no entries, 0 is returned.
- oipmOpenGroup(IN inputDoc, IN name, IN index) - This function looks (within the current group) for an entry array with the given name and an entry with the given index within that array, checks to see if that entry is a group, and if so, changes the "current group" context to that group.
- oipmGet*(IN inputDoc, IN name, IN index, OUT value) - These functions (oipmGetData, oipmGetInt, and oipmGetString) look for an entry array (within the current group) with the given name and a child with the given index and returns the value associated with that child entry.
- oipmCloseGroup(IN inputDoc) - This function switches "current group" to the parent of the current group (the parent of the array within which the current group is located). Equivalent to a "cd .." in a UNIX shell.

[0094] Submodule 64 (core-private) includes a set of implementation "helper" functions that are used by markup implementations to build the parse tree. These helper functions, according to one embodiment, are named: oipmAdd-

Array, oipmiAddEntry, oipmiGetArray, oipmiGetEntry, oipmiNewArray, and oipmiNewEntry, and are characterized as follows:

- oipmiAddArray (IN inputDoc, IN groupEntry, IN entryArray) - adds the given array as a child of the given group entry.
- 5 • oipmiAddEntry (IN inputDoc, IN entryArray, IN entry) - adds the given entry as a child of the given entry array.
- oipmiGetArray (IN inputDoc, IN name, OUT entryArray) - returns the entry array with the given name within the current group.
- oipmiGetEntry (IN inputDoc, IN name, IN index, OUT entry) - returns the entry at position "index" within the array named "name" found within the current group.
- 10 • oipmiNewArray (IN inputDoc, OUT entryArray) - returns a newly initialized entry array
- oipmiNewEntry (IN inputDoc, OUT entry) - returns a newly initialized entry.

[0095] The input module 51 also includes a set of input markup implementations 63, as noted above. Each markup style has an associated input markup implementation. When "oipmInitInputDoc" is called, the markup library calls the appropriate markup implementation, based upon the markup style selected. The markup implementations are expected to parse through the original input buffer (possibly modifying it) and use the helper functions described above to construct the input document parse tree.

[0096] Note that the markup implementations 63 and the implementation helper functions are not visible to users of the WSDK Library 4 -- they are artifacts of markup language implementation.

B. Output Module

[0097] The output module 52 of the markup library 4 is based on a similar principle to that of the input module 51 (i. e., users of the library construct documents via a standardized logical model without regard to markup-language specific issues) but has a different implementation. A strict inversion of the input module would involve creating a parse tree via a set of markup-neutral API functions, followed by a markup-specific translation of that parse tree to a text buffer. For performance reasons, however, only a very sparse amount of context information is maintained and the output document is mostly built on the fly.

[0098] The output module 52 has a structure similar to that of the input module 51 represented in Figure 6. In particular, the output module 52 includes:

- A single complex (core-public-data) data structure, called OipmOutputDoc, used by both markup library users and markup implementations.
- A set of (core-public-func) functions for "writing" the data structure, used by markup library users. For each markup implementation, there is a (markup-private) function that corresponds to (almost every) core-public-func function.
- 35 • A single (core-public-func) function for "initializing" the data structure, used by markup library users. For each markup implementation, there is a (markup-private) function that corresponds to this initialization function. These markup-private functions typically have very little functionality.
- A set of (core-private) functions for "writing" the data structure, used by markup implementations.

[0099] The data structure OipmOutputDoc represents the document under construction. Key fields of this structure include: 1) A markup style, 2) a group "stack"; and 3) the output buffer under construction. The group "stack" is used to capture the context for the output document. For example, suppose we were at the following location in a document:

```

45 <group1>
    <blah>blahblah</blah>
    <group2>
        <blah>blahblah</blah>
        <group3>
            <-- we're here

```

50 [0100] In this case, the group stack would contain:

```

    group1
    group2
    group3 <-- top of stack

```

55 [0101] Similarly to the input module 51, the output module 52 includes a set of key functions for operating on output documents, which according to one embodiment are: oipmInitOutputDoc, oipmBeginGroup, oipmEndGroup, and oipmPut* (where "*" represents one of several possible suffixes). These functions are characterized as follows:

EP 1 345 119 A2

- oipmInitOutputDoc(IN outputDoc, IN markupStyle) - initializes the document appropriately for the given markup style. Note that this function provides much less functionality than its counterpart in the input module 51, oipmInitInputDoc.
- oipmBeginGroup(IN outputDoc, IN name) - open up a group with the given name within the current context (and add the new group name to the top of the group stack). For a markup language based on XML, this would typically involve writing the group name within <>, e.g.:

```
oipmBeginGroup(odoc, "foo")
would add "<foo>" to the output document.
```

- oipmEndGroup(IN outputDoc) - close the current group and remove it from the top of the group stack. For a markup language based upon XML, this would typically involve writing a close tag for the group, e.g.:

```
oipmBeginGroup(odoc, "foo1");
oipmBeginGroup(odoc, "foo2");
oipmEndGroup(odoc);
oipmEndGroup(odoc);
```

might result in:

```
<foo1>
  <foo2>
  </foo2>
</foo1>
```

- oipmPut*(IN outputDoc, IN name, IN value) - write an item of the appropriate type (oipmPutData, oipmPutInt, oipmPutString) into the document, e.g.:

```
oipmBeginGroup(odoc, "group1")
oipmPutString(odoc, "foo", "foo_value");
oipmBeginGroup(odoc, "group2");
oipmPutInt(odoc, "ifoo", 10);
oipmEndGroup(odoc);
oipmEndGroup(odoc);
```

might result in something like:

```
<group1>
  <foo>foo_value</foo>
  <group2>
    <ifoo>10</ifoo>
  </group2>
</group1>
```

[0102] An implementation helper function in the output module 52 reserves space in the output buffer prior to writing to it. This function, oipmiReserveOutput, is defined as follows:

- oipmiReserveOutput(IN outputDoc, IN bytesNeeded)

[0103] The output module 52 also contains a set of output markup implementations, as noted above. Input markup implementations 63 include a single function that is called from oipmInitInputDoc. Output markup implementations, however, include a set of functions to accomplish each of the following actions:

- Initializing the output buffer
- Beginning a group
- Ending a group
- Writing a data buffer
- Writing an integer
- Writing a string
- Finalizing an output buffer

[0104] When the corresponding API function is called by a user of the markup library API, that function in turn calls the appropriate implementation function (based upon the markup style registered for the document).

C Header Files

[0105] In one embodiment of the invention, the above-described data structures and key functions for both the input module 51 and the output module 52 are all declared in a single C header file. The implementation helper functions are declared in a separate header file (for use by markup implementers).

[0106] In one embodiment, these header files are named oip_markup.h, oip_markup_impl.h, oipm_light.h and oipm_soap.h, and are defined as follows:

- oip_markup.h: Contains definitions for OipmInputDoc, OipmOutputDoc, and public oipm* functions. Used by programmers who code to a WSDK-based client API.
- oip_markup_impl.h: Contains declarations for all oipmi* functions. Used by markup implementors.
- oipm_light.h: Contains declarations for the "Lite" markup. Used by markup library.
- oipm_soap.h: Contains declarations for the "SOAP" markup. Used by markup library.

D. Markup Implementations

[0107] In one embodiment of the invention, the WSDK library includes markup implementations to support at least two markups, SOAP and OIP-LIGHT.

1. OIP-LITE Markup Implementation

[0108] The OIP-LITE markup language has a simple syntax. Groups entries are represented as:

```
groupName (
)
```

[0109] Item entries use a simple run-length encoding mechanism for values:

```
FirstName 6
Andrew
```

[0110] (In this case "6" is the length of the value "Andrew" associated with the entry named "FirstName").

[0111] For example:

```
oipmInitOutputDoc(odoc, eOipmLITE);
oipmBeginGroup(odoc, "Person");
oipmPutString(odoc, "FirstName", "Andrew");
oipmPutString(odoc, "LastName", "Large");
oipmPutInt(odoc, "Age", 36);
oipmEndGroup(odoc);
oipmFinalizeOutput(odoc, buf);
```

would result in a buffer that contains:

```

5      Envelope (
        Person (
          FirstName 6
            Andrew
          LastName 5
            Large
10         Age 2
            36
        )
    )

```

```

15  [0112] When parsed into an input document, you would have:
      array "Envelope"
        entry (group) 0
          array "Person"
20            entry (group) 0
              array "FirstName"
                entry (item) 0, value "Andrew"
              array "LastName"
                entry (item) 0, value "Large"
25            array "Age"
              entry (item) 0, value "36"

```

2. SOAP Markup Implementation

30 [0113] The SOAP implementation can be simply an XML parser generator with some hard-coded conventions around special elements, including the SOAP "envelope", "header", and "body" groups. For example:

```

      oipmInitOutputDoc(odoc, eOipmSOAP);
      oipmBeginGroup(odoc, "Person");
35      oipmPutString(odoc, "FirstName", "Andrew");
      oipmPutString(odoc, "LastName", "Large");
      oipmPutInt(odoc, "Age", 36);
      oipmEndGroup(odoc);
      oipmFinalizeOutput(odoc, buf);
40
      would result in:
      <SOAP-ENV:Envelope xmlns:SOAP-
ENV="http://schemas.xmlsoap.org/soap/envelope/">
      <SOAP-ENV:Body>
45      <oip:envelope xmlns:oip="http://www.openwave.com/protocols/oip/">
        <oip:Person>
          <oip:FirstName>Andrew</oip:FirstName>
          <oip:LastName>Large</oip:LastName>
          <oip:Age>10</oip:Age>
50      </oip:Person>
        </oip:envelope>
      </SOAP-ENV:Body>
    </SOAP-ENV:Envelope>

```

and would have exactly the same parse tree structure as the "LITE" version for the input document.

55 [0114] In one embodiment of the invention, the SOAP markup implementation includes the following functions and operates as described:

- oipmsParseDoc() - this is called from oipmInitInputDoc when the markup style is SOAP. Its role is to parse through

EP 1 345 119 A2

a buffer containing SOAP/XML markup and build the corresponding parse tree. Consider the following document:

```
<SOAP-ENV:Envelope ...>
<SOAP-ENV:Body>
  <oip:envelope ...>
5    <oip:AuthUser>
      <oip:userName>user</oip:userName>
      <oip:password>password</oip:password>
    </oip:AuthUser>
  </oip:envelope>
10 </SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

[0115] The function would skip over the SOAP-ENV tags to find the "oip:envelope" : tag, and then:

- 15 • When it found the opening "oip:envelope" tag, it would call:
 - oipmiNewEntry() - to create an entry for the envelope group
 - oipmiNewArray() - to create an array for the envelope group
 - oipmiAddEntry() - to add the envelope entry to the envelope array.
- 20 and then set the newly created entry as the current group.
- When it found the "oip:AuthUser" tag, it would call:
 - 25 • oipmiNewEntry() - to create an entry for the AuthUser group
 - oipmiNewArray() - to create an array for the AuthUser group
 - oipmiAddArray() - to add the AuthUser array to the envelope group
 - oipmiAddEntry() - to add the AuthUser entry to the AuthUser array
- 30 and then set the newly created entry as the current group.
- When it found the "oip:userName" tag, it would call:
 - oipmiNewEntry() - to create an entry for the userName item
 - oipmiNewArray() - to create an array for the userName item
 - 35 • oipmiAddArray() - to add the userName array to the AuthUser group
 - oipmiAddEntry() - to add the userName entry to the username array
- When it found the "user" text for the oip:userName tag, it would change the entry type to ITEM, set the value to point to where "user" occurred in the text buffer, and skip past the "/oip:userName" closing tag.
- 40 • When it found the "oip:password" tag it would call:
 - oipmiNewEntry() - to create an entry for the password item
 - oipmiNewArray() - to create an array for the password item
 - oipmiAddArray() - to add the password array to the AuthUser group
 - 45 • oipmiAddEntry() - to add the password entry to the username array
- When it found the "password" text for the oip:password tag, it would change the entry type to ITEM, set the value to point to where "password" occurred in the text buffer, and skip past the "/oip:password" closing tag.
- When it found the "/oip:AuthUser" closing tag, it would change the current group to the parent of the AuthUser array (the envelope group).
- 50 • When it found the "/oip:envelope" closing tag, it would change the current group to the parent of the envelope array (the top context of the document).
- oipmsBeginGroup - called from oipmBeginGroup when the markup style is SOAP. Example:
 - 55 oipmsBeginGroup(doc, "AuthUser")
 - would call:

```
oipmsReserveOutput(doc, 15);
oipmiOutputData(doc, "<oip:AuthUser>\n");
```

and then push "AuthUser" onto the output group stack.

- 5 • oipmsEndGroup - called from oipmEndGroup when the markup style is SOAP. Example:

```
oipmsEndGroup(doc)
```

would call:

10

```
oipmiReserveOutput(doc, 16);
oipmiOutputData(doc, "</oip:AuthUser>\n");
```

and then pop "AuthUser" off the output group stack.

- 15 • oipmsFinalize - called from oipmFinalizeOutput when the markup style is SOAP. This just adds the SOAP closing tags to the document, e.g.:

```
static const char* soapTrailer = "</SOAP-ENV:Body>\n</SOAP-ENV:Envelope>\n";
oipmiReserveOutput(doc, strlen(soapTrailer));
20 oipmiOutputData(doc, soapTrailer);
```

- oipmsInitDoc - called from oipmInitOutputDoc when the markup style is SOAP. This just adds the SOAP header tags to the document, e.g.:

25

```
static const char* soapHeader = "<SOAP-ENV:Envelope>\n<SOAP-ENV:Body>\n";
oipmiReserveOutput(doc, strlen(soapHeader));
oipmiOutputData(doc, soapHeader);
```

- 30 • oipmsPutData - called from oipmPutData when the markup style is SOAP. This encodes the data in base64 and then writes the data to the document, e.g.:

```
char encodedBuffer[];
base64Encode(buffer, encodedBuffer);
```

35

```
oipmiReserveOutput(doc, strlen(encodedBuffer) + (strlen(name) + 7) * 2);
oipmiOutputData(doc, "<oip:");
oipmiOutputData(doc, name);
oipmiOutputData(doc, ">");
oipmiOutputData(doc, encodedBuffer);
40 oipmiOutputData(doc, "</oip:");
oipmiOutputData(doc, name);
oipmiOutputData(doc, ">");
```

- oipmsPutString - called from oipmPutString when the markup style is SOAP, e.g.:

45

```
oipmiReserveOutput(doc, strlen(buffer) + (strlen(name) + 7) * 2);
oipmiOutputData(doc, "<oip:");
oipmiOutputData(doc, name);
oipmiOutputData(doc, ">");
50 oipmiOutputData(doc, buffer);
oipmiOutputData(doc, "</oip:");
oipmiOutputData(doc, name);
oipmiOutputData(doc, ">");
```

55

VI. OSD and OSDGen

[0116] Development of a WSDK-based service starts with Operation Structure Definition (OSD) specifications for the service. OSD is a hierarchical declaration language with a lexical syntax somewhat similar to LISP. Fundamentally,

OSD is a language for describing the logical structure of operation requests and replies. The structure declared within an OSD file directly maps to markup API calls described above. In other words, OSD is used to describe the interface to a set of operations (similar to the logical section of a WSDL description) and to introduce additional server-side authorization and validation logic.

5 **[0117]** Programmers will manually construct OSD files during development. They will then use OSDGen 1 during development to create skeleton server-side operation implementations (in, for example, C++ or Java), which they then manually populate with the code that handles executing the operation.

[0118] During build (when source code is translated to binary), OSDGen 1 is used to generate the entire service client API, as well as portions of the server-side code (the context classes). The generated code is generated and
10 compiled in a single logical build step.

[0119] OSDGen 1 is not used at runtime. However, both the code generated by OSDGen 1 (the client API and the server-side operation context code) and the manually written server-side operation implementation are used.

[0120] The primary entities defined using OSD include:

- 15 • *service*. This is the primary top-level entity in an OSD file. A *service* has a set of attributes (including name, version number, description, etc.) as well as a number of sub-entities. Sub-entities of *service* include *itemdef*, *groupdef*, *faultdef*, and *operation*.
- 20 • *itemdef*. Itemdefs are used for defining simple types. In addition to a name, they have a corresponding markup type (e.g., string, integer, etc.) and a set of validation rules (min/max value, regular expression, etc.).
- *groupdef*. Groupdefs are used for defining complex "struct"-style types. In addition to a name, they contain a combination of *items* and *groups*:
- 25 • *item*. Items represent simple name/value pairs. In addition to a name, they have an associated *itemdef* or markup type, a cardinality clause, and a default value.
- *group*. Groups represent complex name/value pairs. In addition to a name, they have an associated *groupdef* and a cardinality clause.
- 30 • *faultdef*. Faultdefs are used for defining a service-specific fault that may be returned from an operation defined within the service. Faultdefs contain a name, a numeric code, and a message format string (similar to a *printf* string, but with argument names as well).
- 35 • *operation*. Operations represent the individual tasks that a service can execute. Operations have a set of attributes (including name, version number, description, etc.) as well as sub-entities.
- *request*. A request defines the structure of the information provided in a request for execution. Structurally, it is equivalent to a *groupdef* with no name.
- 40 • *reply*. A reply defines the structure of the information provided in a reply provided with an operation execution. Structurally, it is equivalent to a *groupdef* with no name.
- *fault*. This is used to identify that a particular fault (identified by *faultdef* name) may be returned by this
45 operation.

[0121] OSD also supports the equivalent of the C pre-processor "#include" directive.

[0122] The most common organization of OSD files includes: a single top-level service file; a file that contains all the *itemdef*, *groupdef*, and *faultdef* definitions for the service; and a separate file for each operation.

50 **[0123]** Once OSD files exist, OSDGen 1 can be used to generate the following types of service-specific source files (based upon command-line options to OSDGen 1), server-side operation classes, server-side operation context classes, server-side service facility classes, client-side operation context code, and client-side support files:

- 55 • Server-side operation classes. For each defined operation that does not have an associated C++ class, one will be generated for it. The generated classes inherit from the *OipsOperation* class and contain a method called *execute()*. This method is invoked (with appropriate parameters) when execution is requested for that operation.
- Server-side operation context classes. For each defined operation, a C++ context class will be generated. Oper-

ation context classes inherit from OipsOpCtxt and contain request and reply data for a particular instance of an operation execution and logic for serializing/deserializing operation request and reply data.

- Server-side service facility classes. For each service, a small set of additional C++ classes are generated. These include a top-level class derived from OipsService (includes service-initialization and operation registration logic) as well as classes that define service-specific types and itemdef-specific validation logic.
- Client-side operation context code. For each operation, a "struct" declaration and associated functions is generated. This code is designed to plug into the WCAF 2 and is similar to the server-side operation context classes, with the exception that it lacks the rich validation and version mapping logic found in the server-side classes and is written in C instead of C++.
- Client-side support files. These files include both WSDL specifications for the service and fault message catalogs (in *catgets* format, to enable localization of fault messages).

2. OSD Details

[0124] OSD has a grammar that maps to the markup language structure. There are two types of "entries" in OSD: a "group" entry and an "item" entry. A group entry in OSD is depicted as:

```
(type "name"...)
```

[0125] An "item entry" in OSD is depicted as:

```
name "value"
```

[0126] The following main groups are defined within OSD: service, itemdef, item, groupdef, group, operation, which are characterized as follows:

- service - used to define services. This is the highest-level construct in OSD. Everything else falls within a service definition, e.g.:

```
(service "smg"
```

```
)
```

- itemdef - used to define types for items that will be contained in a request or reply, e.g.:

```
(itemdef "OipAuthToken"
  mtype "string" # markup type (string, int, or data)
  min_length "64"
  max_length "64")
```

- item - an instance of an item, e.g.:

```
(item "authToken"
  type "OipAuthToken"
  card "1..1")
```

- groupdef - used to define (complex) types for groups that will be contained in a request or reply, e.g.:

```

(groupdef "Subscriber"
  (item "id" type "SmsgSubscriberId" card "1..1")
  (item "authToken" type "OipAuthToken" card "1..1")
  (group "attrs" type "Attribute" card "0..n")
)

```

- group - used to define an instance of a group
- operation - used to define an operation. Every operation must contain a "request" and a "reply" group:

```

(operation "AuthenticateSubscriber"
  (group "request"
    (group "lookupInfo"
      type "LookupInfo"
      desc "Information to find the subscriber."
      card "1..1")
    (group "authInfo"
      type "AuthInfo"
      desc "Information to authenticate the subscriber."
      card "1..1")
    )
  (group "reply"
    (item "id"
      type "SmsgSubscriberID"
      card "1..1"
      write_ctxt "true"
      desc "Globally unique subscriber ID")
    (item "authToken"
      type "OipAuthToken"
      card "1..1"
      write_ctxt "true"
      desc "To be used with future requests.")
    )
  )
)

```

[0127] Consider the following trivial example of a service that contains an operation for translating from a numerical digit to the corresponding English word (e.g., "1" -> "one").

```

5      (service "digit"
      (itemdef "DigitWord"
        mtype "string"
        min_length "3"
        max_length "5")

10     (itemdef "DigitSymbol"
        mtype "integer"
        min "0" max "9")

15     (operation "ToWord"
        (group "request"
          (item "symbol" type "DigitSymbol" card "1..1"))
        (group "reply"
          (item "word" type "DigitWord" card "1..1"))

20    )
  )

```

25

[0128] On the client side, OSDGen 1 would produce the following set of C declarations:

```

30      struct DigitToWordRequest {
        int symbol;
      };

      struct DigitToWordReply {
35        char* word;
      };

      struct DigitToWordOperation {
40        OipcStatus (*requestToDoc)(
          OipcCall* call,
          const OipcOperation* op);

45        OipcStatus (*docToReply)(
          OipcCall* call,
          OipcOperation* op);

50        void (*cleanup)(
          OipcOperation* op);

      DigitToWordRequest request;
55      DigitToWordReply reply;
    };

```

extern void digitInitToWord(DigitToWordOperation*);

[0129] The following is an example of what the init and "requestToDoc" functions for DigitToWordOperation would look like:

```

5      void digitInitToWord(
        DigitToWordOperation* op)
        {
10         memset(op, 0, sizeof(DigitToWordOperation));
            op->cleanup = _DigitToWordCleanup;
            op->docToReply = _DigitToWordDocToReply;
            op->requestToDoc = _DigitToWordRequestToDoc;
15         }

        static OipsStatus _DigitToWordRequestToDoc(
            OipcCall* call,
20             const OipcOperation* op)
            {
                const DigitToWordOperation* l_op = (const DigitToWordOperation*)op;
                const DigitToWordRequest* r = &l_op->request;
25

                oipmBeginGroup(&call->requestDoc, "digit_ToWord");
                oipmPutInt(&call->requestDoc, "symbol", r->symbol);
                oipmEndGroup(&call->requestDoc);
30            }

```

Server-side context:

35 **[0130]** If a reply item is marked with 'write_ctxt "true"', then the server will write the reply data for that item in two locations within the OipsCallCtxt object: 1) in the reply document (using the same oipm* functions displayed above for the client); and 2) in a context "map" maintained within the call context (see below). If a request item is marked with 'read_ctxt "true"', then a query for that item in the request will look in two places: 1) in the request document (via the appropriate oipmGet* function); and 2) if it is not found in the request document, then the context map is queried for the item.

40 **[0131]** This relationship between write_ctxt and read_ctxt is an important facet of the batching capabilities of the server, since it allows for a batch in circumstances that might normally require an additional round-trip. In particular, one can specify that reply (output) data from one operation in a batch be used as request data (input) for a later operation in the same batch. For example, consider the following:

45 Operation AuthUser:

 request = user name and password
 reply = authentication token

50

 Operation GetUserProfile:

 request = authentication token
 reply = profile

55

[0132] Marking the authentication token in the reply of AuthUser as 'write_ctxt "true"' and the authentication token in the request of GetUserProfile allows for both of these operations to be submitted as part of the same batch (essentially using output from AuthUser as input to GetUserProfile).

[0133] As an example of how OSDGen 1 parses OSD files, consider the following simple example OSD code fragment:

```

5      (groupdef "Person"
      (item "FirstName" mtype "string")
      (item "LastName" mtype "string")
      (item "Age" mtype "int" min "1" max "200")
10     )

      (operation "PromoteEmployee"
      (request
15         (group "person" type "Person"))
      ...
      )

```

[0134] The above OSD fragment says that a "PromoteEmployee" operation expects to have a "Person" identified in the request. OSDGen 1 parses these structural definitions and automatically generates the code that actually invokes the WSDK library API. So, for the above operation, client code would be automatically generated by OSDGen 1 to contain the following code fragment in the portion of the "PromoteEmployee" client API that is responsible for generating the request message:

```

25     ...
      oipmBeginGroup(odoc, "Person");
      oipmPutString(odoc, "FirstName", request.FirstName);
30     oipmPutString(odoc, "LastName", request.LastName);
      oipmPutString(odoc, "Age", request.Age);
      oipmEndGroup(odoc);
35     ...

```

VII. System Operation

A. System Initialization

[0135] System initialization begins with the initialization of server-level resources (logging, configuration, etc.) for a server instance, followed by discovery and dynamic loading of the set of services to be hosted within that server instance (during initialization, each service creates and registers a single instance of each operation class). The call manager, oipsCallMgr, is then initialized and created.

B. Document Exchange

[0136] An example of a document exchange between client and server using the WSDK APIs is as follows, as illustrated in Figure 7:

1. Client initializes an output document (block 701) by calling oipmInitOutputDoc, selecting the markup as appropriate for its performance and compatibility needs. For example, if a large volume of non-text data (e.g., .WAV or MP3 files) is to be sent, then the client might opt to use a markup other than SOAP, which is not well-suited for such uses.
2. (Assuming use of the above APIs) Client uses the WSDK markup library 4 to add groups and items to the document as appropriate for the services being requested and then asks for the "serialized" representation of the output document (a data buffer containing the "on the wire" representation of the document) (block 702). This is done by making a number of calls to oipmPut*(), oipmBeginGroup() and oipmEndGroup(), as appropriate, followed

by a call to `oipmFinalizeOutput()` (which allows the implementation to append any markup-specific trailer onto the buffer and returns the buffer).

3. Client transmits the document to the server (block 703), identifying the markup style in a transport-specific manner. For a typical case where the transport is via an HTTP POST, a URI is used to identify the markup style (e.g., URL is "http: / /host:port/oip/soap" or "http://host:port/oip/light").

4. Server reads the request message from client into a data buffer (from socket, file, Apache module API, etc.) (block 704).

5. Server invokes the markup library 4 to "parse" the document into the input structure (passing the appropriate markup style to the markup library 4) and initializes an empty output document in the same markup style in a data buffer (block 705). In this context, "parsing" includes: 1) parsing from text form into "input document" form, and 2) extracting data from the input document to build a native language structure.

6. Server invokes the `handleCall` method of `oipsCallMgr` to execute the instructions contained in the input document, to populate the output document (block 706).

7. Server transmits the output document back to the client, for example as the reply to the HTTP POST (block 707).

8. Client receives and parses the document into an input structure (using the markup library 4), passing in the original requested markup style (block 708).

9. Client traverses the document to determine the results of its request (block 709).

B. Client Operation

[0137] The actions performed by the client during a document exchange will now be described in greater detail, with reference to Figure 8. The actions performed by the client are as follows :

1. Create and initialize a "call" object, specifying the markup style and error handling mode desired (block 801).

2. For each operation to bundle into the call (blocks 802, 806, 812):

a. Create and initialize an appropriately typed Operation structure (block 803).

b. Populate the "request" sub-structure (block 804).

c. Add the operation to the call (block 805).

3. Invoke the "Generate Request" method of the call object (call `oipcGenerateRequest`) to obtain the markup-specific serialized representation of the request (block 807).

4. Send the request message to the server using an appropriate transport engine and obtain the reply message (block 808).

5. Invoke the "Set Reply" method of the call object (call `oipcSetReply`) to parse the reply message (block 809).

6. Examine the "reply" sub-structures as appropriate for each of the operations in the call (block 810).

7. Clean-up the call object (block 811).

C. Server Operation

[0138] As discussed above, `OipsCallMgr` is the top-level manager of the call dispatch process on the server side. `OipsCallMgr` has a "handleCall" method, which is invoked by an underlying server framework. The `handleCall` method performs the following process, illustrated in Figure 9, when a new request message is received:

1. Construct a call context, `OipsCallCtxt`, passing in the request message and the markup style (block 901).

2. If the error mode is "rollback" (block 902), initialize the transaction context (block 903).

3. Call `oipmNumEntries()` on the `OipsCallCtxt`'s request document to find out how many operations exist in the document (block 904).

4. For each operation in the document (blocks 905-916):

a. If the error mode is not "rollback" (block 906), initialize the transaction context (block 907).

b. Request construction of the appropriate operation context object, `OpCtxt`, from the service manager and add it to the transaction manager (block 908).

- Service manager reads the service name from the request and asks the corresponding service for the context.

- The service reads the operation name from the request and asks the corresponding operation object

for the context.

- The operation object constructs and returns an appropriately typed operation context object.

5 c. Call the operation context's parseRequest() method (block 909).

- The operation context object uses OipsCallCtxt to find and validate the appropriate groups and items and populates "request" structure (data is not copied into request structure - only pointers).

10 d. Find the appropriate Operation object and call its "execute()" method, passing the OpCtxt and transaction manager objects to it (block 910).

- e. If the operation execute() succeeds (block 911), and if the error mode is not "rollback" (block 912), invoke "commit()" on the transaction manager (which will write the reply and invoke commit on anything the operation may have added to the transaction manager) and destroy the transaction manager (block 913).

15 f. If the operation execute() fails (block 911), then:

- Call "rollback()" on the transaction manager (block 914). This will call "rollback()" on anything contained within the transaction manager (and will write error messages as the replies for any OpCtxt objects).
- If the error mode is not "continue" (block 915), stop execution (block 916) and return (block 919).

20 5. If the error mode is "rollback" (block 917), call "commit()" on the transaction manager (see 4e) (block 918).

VIII. Processing System Implementation

25 **[0139]** The techniques described above may be implemented in one or more processing systems using software executed by one or more conventional processors, or using special-purpose hardwired circuitry, or a combination thereof. Figure 10 is a high-level block diagram showing an example of a processing system that can be used to implement the described techniques. Note that a variety of specific architectures can be represented by Figure 10.

30 **[0140]** The illustrated system includes one or more processors 101, i.e. a central processing unit (CPU), random access memory (RAM) 102, read-only memory (ROM) 103, and, which may be coupled to each other by a bus system 109. The processor(s) 101 may be, or may include, one or more programmable general-purpose or special-purpose microprocessors, digital signal processors (DSPs), programmable controllers, application specific integrated circuits (ASICs), programmable logic devices (PLDs), or a combination of such devices.

35 **[0141]** The bus system 109 includes one or more buses or other connections, which may be connected to each other through various bridges, controllers and/or adapters, such as are well-known in the art. For example, the bus system 109 may include a "system bus", which may be connected through one or more adapters to one or more expansion buses, such as a Peripheral Component Interconnect (PCI) bus, HyperTransport or industry standard architecture (ISA) bus, small computer system interface (SCSI) bus, universal serial bus (USB), or Institute of Electrical and Electronics Engineers (IEEE) standard 1394 bus (sometimes referred to as "Firewire").

40 **[0142]** Also coupled to the bus system 109 are one or more data communication devices 104, one or more mass storage devices 105, one or more input devices 106 and one or more output devices 1010. The input devices 106 may include, for example, a keyboard and a pointing device. The output devices 1010 may include, for example, a display device and an audio speaker. Of course, such devices are unnecessary on a server if the server is not intended to directly interface with a user or operator. It will be understood that the processing system may also include other conventional components that are not shown.

45 **[0143]** Each data communication device 44 is a device suitable for enabling the processing system to communicate with remote devices and may be, for example, a conventional modem, a Digital Subscriber Line (DSL) modem, a cable modem, an Ethernet adapter, an Integrated Services Digital Network (ISDN) adapter, a satellite transceiver, a wireless transceiver (e.g., in the case of a mobile device), or the like. Each mass storage device 105 may include any one or more storage media suitable for storing large volumes of data in a non-volatile manner, such as a magnetic disk or tape, magneto-optical (MO) storage device, or any of various forms of Digital Versatile Disk (DVD) or Compact Disk (CD) based storage, or a combination thereof.

50 **[0144]** The processes described above may be implemented using a combination of instructions and data 108 stored in a storage medium, which may be, for example, any of RAM 102, mass storage device 105 and/or ROM 103, as shown, or on a remote processing system.

55 **[0145]** Thus, a method and apparatus for requesting and performing batched operations for Web Services have been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be recognized that the invention is not limited to the embodiments described, but can be practiced with modification

and alteration within the spirit and scope of the appended claims. Accordingly, the specification and drawings are to be regarded in an illustrative sense rather than a restrictive sense.

5 Claims

1. A method comprising:

10 dynamically determining in a client a plurality of batched operations to be performed by a server for the client during execution of a web based network application, by selecting a portion of an operation structure previously defined using a structured language;
operating the client to specify the plurality of batched operations in a request; and
15 sending the request from the client to the server to request the server to perform the plurality of operations in a batched mode.

2. A method as recited in claim 1, wherein the plurality of batched operations are determined by the client dynamically based on a current state of execution of the application.

3. A method as recited in claim 1, wherein one of the plurality of batched operations uses output of another one of the plurality of batched operations as input, when the plurality of batched operations are executed.

4. A method as recited in claim 1, wherein the request further specifies an error handling mode of a plurality of selectable error handling modes to determine how an error that occurs during execution of the plurality of operations is to be handled, the error-handling mode having been specified by the client.

5. A method as recited in claim 4, wherein the plurality of selectable error handling modes include:

25 a stop mode in which execution of the operations is stopped in response to an error occurring during execution of the operations;
30 a continue mode in which execution of the operations continues in response to an error occurring during execution of the operations; and
a rollback mode in which execution of the operations is stopped and a state that existed prior to execution of the operations is resumed, in response to an error occurring during execution of the operations.

6. A method as recited in claim 1, wherein the client is a mobile client operating on a wireless network and the server operates on a wireline network.

7. A method as recited in claim 1, wherein the request is embodied in a markup language document.

8. A method for a server to provide a service to a client during execution of a web based application in a network environment, the method comprising:

generating a request by the client for the server to perform the service, including
45 creating a call structure including an array specifying a plurality of batched operations to be performed by the server, by selecting a portion of an operation structure previously defined using a structured language, wherein the plurality of batched operations are determined by the client dynamically based on a current state of execution of the application, and wherein the call structure further includes an indication of a markup language used by the client and an error handling mode of a plurality of selectable error handling modes to determine how an error that occurs during execution of the plurality of operations is to be handled, and
50 including the call structure in a markup language document; and
sending the call structure in a markup language document from the client to the server to request the service, such that the server performs the plurality of operations in the batched mode.

9. A method as recited in claim 8, wherein the plurality of selectable error handling modes include:

55 a stop mode in which execution of the operations is stopped in response to an error occurring during execution of the operations;
a continue mode in which execution of the operations continues in response to an error occurring during

execution of the operations; and
a rollback mode in which execution of the operations is stopped and a state that existed prior to execution of the operations is resumed, in response to an error occurring during execution of the operations.

- 5 10. A method as recited in claim 8, wherein the client is a mobile client operating on a wireless network and the server operates on a wireline network.
11. A method as recited in claim 8, wherein one of the plurality of batched operations uses output of another one of the plurality of batched operations as input, when the plurality of batched operations are executed.
- 10 12. A method comprising:

receiving a request at a server from a client during execution of a web based network application, the request specifying a plurality of batched operations associated with a service to be provided by the server to the client,
15 the plurality of batched operations having been determined by the client dynamically by the client selecting a portion of an operation structure previously defined using a structured language; and
operating the server to perform the plurality of operations in a batched mode.
13. A method as recited in claim 12, wherein the plurality of batched operations have been determined by the client based on a current state of execution of the application.
14. A method as recited in claim 12, wherein one of the plurality of batched operations uses output of another one of the plurality of batched operations as input, when the plurality of batched operations are executed.
- 25 15. A method as recited in claim 12, wherein the request further specifies an error handling mode of a plurality of selectable error handling modes to determine how an error that occurs during execution of the plurality of operations is to be handled, the error-handling mode having been specified by the client.
16. A method as recited in claim 15, wherein the plurality of selectable error handling modes include:
30 a stop mode in which execution of the operations is stopped in response to an error occurring during execution of the operations;
a continue mode in which execution of the operations continues in response to an error occurring during execution of the operations; and
35 a rollback mode in which execution of the operations is stopped and a state that existed prior to execution of the operations is resumed, in response to an error occurring during execution of the operations.
17. A method as recited in claim 12, wherein the client is a mobile client operating on a wireless network and the server operates on a wireline network.
- 40 18. A method as recited in claim 12, wherein the request is embodied in a markup language document.





FIG. 1

**FIG. 2**




FIG. 3




FIG. 4A




FIG. 4B




FIG. 5





FIG. 6

**FIG. 7**






FIG. 8

**FIG. 9A**

**FIG. 9B**

